

Claim(s)

1. A method of storing data in a cache memory of a storage device, comprising:
  - providing access to a first segment of the cache memory on behalf of a first group of external host systems coupled to the storage device; and
  - 5        providing access to a second segment of the cache memory on behalf of a second group of external host systems coupled to the storage device, wherein at least a portion of the second segment of the cache memory is not part of the first segment of the cache memory.
2. A method, according to claim 1, wherein no portion of the second segment of the cache memory is part of the first segment.
- 10    3. A method, according to claim 1, further comprising:
  - providing a first data structure in the first segment of the cache memory; and
  - providing a second data structure in the second segment of the cache memory, wherein
  - accessing the first segment includes accessing the first data structure and accessing the second segment includes accessing the second data structure.
- 15    4. A method, according to claim 3, wherein the data structures are doubly linked ring lists of blocks of data.
5. A method, according to claim 4, wherein each block of data corresponds to a track on a disk drive.

6. A method, according to claim 1, further comprising:
- apportioning the cache memory into slots; and
- mapping each of the slots to at least one of the first and second segments of the cache memory.
- 5 7. A method, according to claim 6, wherein the slots are mapped to the segments using a formula.
8. A method, according to claim 6, wherein the slots are mapped to the segments using a table.
9. A method, according to claim 1, wherein the groups are mapped to particular ones of the segments using a table.
- 10 10. A method, according to claim 9, wherein the table includes group identifiers and corresponding masks.
11. A method, according to claim 10, wherein the masks are binary values that have a "one" bit in an Nth bit position to indicate that a group is assigned to an Nth segment.
12. A method, according to claim 1, further comprising:
- 15 in response to a request for a block of cache memory, determining availability of a block of cache memory for a group mapped to the cache memory.

13. A method, according to claim 12, further comprising:

in response to no blocks of cache memory for the group being available, providing a block of cache memory corresponding to another group.

14. A method, according to claim 13, wherein the block of cache memory that is provided is at

5 least one of: a next available block, a block corresponding to a group having a greatest number of blocks assigned thereto, a block corresponding to a group having a greatest number of available blocks, and a block corresponding to a group having a greatest percentage of available blocks.

15. A cache memory of a storage device, comprising:

10 a first segment of the cache memory that is accessed on behalf of a first group of external host systems coupled to the storage device; and

a second segment of the cache memory that is accessed on behalf of a second group of external host systems coupled to the storage device, wherein at least a portion of the second segment of the cache memory is not part of the first segment of the cache memory.

16. A cache memory, according to claim 15, wherein no portion of the second segment of the

15 cache memory is part of the first segment.

17. A cache memory, according to claim 15, further comprising:

a first data structure in the first segment of the cache memory; and

a second data structure in the second segment of the cache memory, wherein accessing the first segment includes accessing the first data structure and accessing the second segment

5 includes accessing the second data structure.

18. A cache memory, according to claim 17, wherein the data structures are doubly linked ring lists of blocks of data.

19. A cache memory, according to claim 18, wherein each block of data corresponds to a track on a disk drive.

10 20. A cache memory, according to claim 15, further comprising:

a plurality of slots, each corresponding to a portion of the cache memory, wherein each of the slots is mapped to at least one of the first and second segments of the cache memory.

21. A cache memory, according to claim 20, wherein the slots are mapped to the segments using a formula.

15 22. A cache memory, according to claim 20, wherein the slots are mapped to the segments using a table.

23. A cache memory, according to claim 15, wherein the groups are mapped to particular ones of the segments using a table.

24. A cache memory, according to claim 23, wherein the table includes group identifiers and corresponding masks.

5 25. A cache memory, according to claim 24, wherein the masks are binary values that have a "one" bit in an Nth bit position to indicate that a group is assigned to an Nth segment.

26. A storage device, comprising:

a plurality of disk drives;

a plurality of disk interface units, each being coupled to one of said disk drives;

10 a bus that interconnects said disk interface units; and

a cache memory, coupled to said bus, said cache memory having a first segment that is accessed on behalf of a first group of external host systems coupled to the storage device and a second segment that is accessed on behalf of a second group of external host systems coupled to the storage device, wherein at least a portion of the second segment of the cache memory is not  
15 part of the first segment of the cache memory.

27. A method of storing data in a cache memory of a storage device, comprising:

providing access to a first segment of the cache memory on behalf of a first group of external host systems coupled to the storage device;

5 providing access to a second segment of the cache memory on behalf of a second group of external host systems coupled to the storage device, wherein at least a portion of the second segment of the cache memory is not part of the first segment of the cache memory;

in response to a request for a block of cache memory by an external host system of the first group, determining availability of a block of cache memory in the first segment of the cache memory; and

10 in response to no blocks of cache memory in the first segment being available, providing a block of cache memory from the second segment for use by the external host system of the first group.

28. A method, according to claim 27, further comprising:

providing a first data structure in the first segment of the cache memory; and

15 providing a second data structure in the second segment of the cache memory, wherein accessing the first segment includes accessing the first data structure and accessing the second segment includes accessing the second data structure.

29. A method, according to claim 28, wherein the data structures are doubly linked ring lists of blocks of data.

20 30. A method, according to claim 29, wherein each block of data corresponds to a track on a disk drive.